

02-6220

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

ALTERNATE NON-VOLATILE MEMORY FOR ROBUST I/O

BY

CHRISTOPHER MCCARTY
4920 KASHMIRE
COLORADO SPRINGS, CO 80920
CITIZEN OF USA

JEFFREY ROGERS
2520 LYNCREST DR
COLORADO SPRINGS, CO 80918
CITIZEN OF USA

BRUCE TRUNCK
1230 FAWNWOOD RD
MONUMENT, CO 80132
CITIZEN OF USA

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number EV 303 409 336 US

Date of Deposit: December 8, 2003

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

BY:


ReNea D. Berggren

ALTERNATE NON-VOLATILE MEMORY FOR ROBUST I/O

FIELD OF THE INVENTION

[0001] The present invention generally relates to the field of memory devices, and particularly to an internal memory for a bus interface controller card or integrated circuit chip.

BACKGROUND OF THE INVENTION

[0002] Data availability and data protection are key aspects of storage solutions and determine the reliability of a system. The more robust an algorithm is in providing data availability and data protection, the more reliability the system has and the more the user's information will be protected. Users can typically tolerate short periods of time when they cannot access the information in the storage subsystem, but they cannot tolerate the permanent loss of their data. There are many reasons that data can get lost in the storage subsystem. The following is a short list of examples: hardware failure, software failure, and power failure. Some of these can be solved with the following approaches. One way to protect against hardware failure is with the use of RAID algorithms. These algorithms use redundant information on other drives to recreate the data on the failed drive when a new drive replaces the failed drive. These algorithms are standard practice in the industry. But, RAID only protects against a hardware failure. There are still potential ways that data can be lost when power is lost during the write of data to the drive. These writes take a finite amount of time and if the power is lost during this time, then the data that is being written may only get partially written or not at all. This can have a large impact, not only on the integrity of the data, but also on the consistency of the RAID algorithm. The partially written data will no longer match the redundant data on another drive. RAID controllers remove the data integrity problem by adding a battery-backed cache to the RAID controller. The cache is used as a buffer to store the write data until it has been successfully written to disk. Once that data has been successfully written to disk, the data can be deleted from the cache. The battery is used

to keep the write data valid across power cycles. After power is lost and the data has not been successfully written to disk, then the controller will reattempt to write the data to the disk after the power has returned. The disadvantage of this solution is the cost of the battery and the cost of the external memory for the cache data.

[0003] One of the biggest issues with the RAID algorithms is to maintain the consistency of the data between the real data and the redundant data information. This is also known as keeping the data synchronized. One of the ways that data can get unsynchronized is when the write to the real data disk completes, but the write to the redundant data information disk does not complete successfully. The solution to this problem is to detect when one write completes and the other write does not. Logging the write information to each disk and keeping track of when each one completes can detect when the writes do not complete. This is known as Write Journaling. The logging information must be kept in some type of non-volatile memory to keep the logged information valid across power cycles. The entry in the write journal will be deleted when all writes complete. If an entry is detected at power-on, then the data can be resynchronized only on the region affected by this write command. Placing the non-volatile memory external to the controller is not a cost effective solution.

[0004] Therefore, it would be desirable to provide a method, circuit, and system for providing an internal memory to a bus interface controller card or chip.

SUMMARY OF THE INVENTION

[0005] Accordingly, the present invention is directed to a method, circuit, and system for providing an internal memory to a bus interface controller card or chip by providing a back up power source in case general power to the card or chip fails.

[0006] The present invention offers a cost effective solution is to bring the non-volatile memory internal to the controller, thus eliminating the external part. Accordingly, memory is brought internal to the controller. This can be provided in a variety of ways depending on how long the memory is needed to be non-volatile. An external battery provides the non-volatile aspect of the internal memory in a cost effective way while retaining board functionality.

[0007] It is to be understood that both the forgoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

FIG. 1 illustrates an embodiment of the general system of the present invention;

FIG. 2 illustrates an embodiment of a specific system of the present invention;

FIG. 3 illustrates an embodiment of a method for creating a system or circuit of the present invention;

FIG. 4 illustrates a method of applying power to the internal memory in the present invention; and

FIG. 5 illustrates an embodiment of a write journaling method of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0009] Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.

[0010] The present invention relates to a method and circuit for locating a non-volatile memory internal to the controller, especially for a Small Computer System Interface (SCSI) controller. For memory to be non-volatile, it must maintain its data while the system power has been removed. This can be achieved using internal volatile memory and an external battery to power the internal memory during the system power down time. The battery is usually much more cost effective than an external non-volatile memory part, thus lowering the overall solution cost of a write journaling implementation. The type of write journaling addressed by the present invention is at the input/output (IO) level; especially, for the Small Computer System Interface (SCSI) CDB level. In IO write journaling, the operating system (OS) file system driver sends IOs (e.g., SCSI CDBs) to the SCSI controller. The firmware (FW) inside the SCSI controller logs these IOs to the journal. After the IO is completed, the FW will remove the logged entry from the journal. By logging the IOs in the journal, it can be determined whether all the IOs have been completed, even when power is interrupted or a system reset occurs. If IOs are detected in the journal at power up or system reset, then the FW can take action to keep the RAID volume consistent (synchronized).

[0011] FIGS 1 and 2 illustrate block diagrams of a general embodiment and a more specific embodiment of the present invention. In FIG. 1, a bus controller interface card or integrated circuit chip 20 contains an internal volatile memory 10, such as a random access memory (RAM). The internal volatile memory is provided with power from an external power source 30. The external power source may be an alternating current to

direct current (AC-to-DC) converter which receives 110 volt, 60 Hertz power and transforms it to a DC voltage through a bridge, transformer, smoothing capacitors, electrolytic capacitors, filters, and/or the like. In the preferred embodiment, as shown in FIG. 2, the power source is a battery 32. The battery life is one of the key parameters with this solution. Batteries only have a limited life while they are being used, so limiting the time when the battery is in use will prolong the life of the battery. The following lists two ways to prolong the life of the battery: 1) Only use the battery power when the system is off. Requires a way to switch from system power to battery power. 2) Only use the battery power when there is data in the Write Journal and the system power is off.

[0012] FIG. 3 illustrates an embodiment of a method for creating a system or circuit of the present invention. An internal memory is provided on a bus interface board or integrated circuit chip 50. The internal memory may be located on a bus interface integrated circuit chip or may be a separate integrated circuit chip of the bus interface controller board. An external power supply is provided to the internal memory 60. Preferably, the external power supply is switched on only when power is not being supplied to the bus interface controller board.

[0013] FIG. 4 illustrates a method of applying power to the internal memory in the present invention. The bus interface controller board is initially powered up 110. A determination is made as to whether power to the bus interface controller board has been interrupted 120. The interruption may be a power down condition or a power instability condition. If power has not been interrupted, then the bus interface board power continues to be used to power the internal memory of the bus interface controller board 150. Otherwise, the external power source is used to power internal memory 130. A determination is made as to whether bus interface controller board power has been restored 140. If it has not, processing returns to step 130. If it has, switching circuitry switches to power the internal memory from the bus interface controller board 150.

[0014] FIG. 5 illustrates an embodiment of a write journaling method of the present invention. After initialization, write information to the real data disk and the redundant information disk is logged in the internal memory 210. If it is determined that all writes have been completed 220, all log entries are deleted from the internal memory 250. Otherwise, a determination is made as to whether power has been removed 230. If it has not, process returns to step 220. That is, write journaling continues. If it is determined that power has been removed or interrupted significantly 230, then a determination is made as to whether power has been restored 240. As long as it is determined that power is removed 240, a wait loop is in effect. When it is determined that power is restored 240, incomplete write commands are processed. Incomplete write commands will only be processed if power is interrupted and entries remain in the write journal when power is reapplied. In an alternative embodiment, instead of deleting log entries from the internal memory, the log entries may be transferred from the internal memory to a mass storage device or other non-volatile memory.

[0015] The present invention may be practiced in a variety of other ways. The bus interface controller board or integrated circuit chip may include a Peripheral Component Interconnect (PCI) controller, a Small Computer System Interface (SCSI) controller, a Redundant Array of Independent Disks (RAID) controller, or another type of controller.

[0016] It is believed that the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages, the form hereinbefore described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.